

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

As rescanning documents *will not* correct images,
please do not report the images to the
Image Problem Mailbox.



UNITED STATES PATENT APPLICATION

FOR

METHOD AND APPARATUS FOR PERFORMING
LOGIC REPLICATION IN FIELD PROGRAMMABLE GATE ARRAYS

Client Docket Nos.: A1258
Attorney Docket No.: ALT.P023

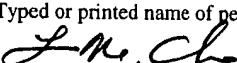
Inventors: Deshanand Singh
Gabriel Quan
Terry Borer
Valavan Manohararajah
Paul McHardy
Ivan Hamer
Karl Schabas
Kevin Chan

Filed By:
Lawrence M. Cho
P.O. Box 2144
Champaign, IL 61825
(217) 377-2500

EXPRESS MAIL CERTIFICATE OF MAILING
"Express Mail" mailing label number ER873666352US
Date of Deposit 3/23/2004

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to the Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450

Lawrence M. Cho
(Typed or printed name of person mailing paper or fee)


(Signature of person mailing paper or fee)

METHOD AND APPARATUS FOR PERFORMING
LOGIC REPLICATION IN FIELD PROGRAMMABLE GATE ARRAYS

RELATED APPLICATION

5 This application is a continuation-in-part, and claims the benefit of co-pending U.S. Application No. 10/617,502 filed on July 11, 2003 entitled "Method and Apparatus for Performing Incremental Placement for Layout-Driven Optimizations on Field Programmable Gate Arrays" under 35 U.S.C., §120, and U.S. Provisional Application 60/486,619 filed on July 11, 2003 entitled "Method and Apparatus for Performing Logic Replication in Field

10 Programmable Gate Arrays" under 35 U.S.C. §119(e).

FIELD OF THE INVENTION

The present invention relates to the field of field programmable gate arrays (FPGAs). More specifically, the present invention relates to a method and apparatus for performing targeted logic replication for timing driven FPGAs using tools such as electronic design automation (EDA) tools.

BACKGROUND

FPGAs may be used to implement large systems that include millions of gates and megabits of embedded memory. Of the tasks required in managing and optimizing a design, placement of components on the FPGAs and routing connections between components on the FPGA utilizing available resources can be the most challenging and time consuming. In order to satisfy placement and timing specifications, several iterations are often required to determine how components are to be placed on the target device and which routing resources to allocate to the components. The complexity of large systems often requires the use of EDA tools to manage and optimize their design onto physical target devices. Automated placement and routing algorithms

in EDA tools perform the time consuming task of placement and routing of components onto physical devices.

The design of a system is often impacted by the connection delays routed along the programmable interconnect of the target device. The interconnect provides the ability to implement arbitrary connections, however, it includes both highly capacitive and resistive elements. The delay experienced by a connection is affected by the number of routing elements used to route the connection. Traditional approaches for reducing the delay were targeted at improving the automated placement algorithms in the EDA tools. Although some reductions in delay were achieved with these approaches, the approaches were not able to perform further improvements to the system after the placement phase. It is often only after the placement phase of the FPGA computer automated design (CAD) flow when connection delays are fully known.

Thus, what is needed is an efficient method and apparatus for performing layout-driven optimizations on FPGAs after the placement phase of the FPGA CAD flow.

SUMMARY

According to an embodiment of the present invention, logic replication is utilized to generate a duplicate of a component. The duplicate is used to reduce the length of a critical path associated with the component without significantly affecting the length of other paths in the system. Logic replication includes making an identical copy of a component, which may include a combination of a LUT and a register, on a critical or near-critical path. The duplicate is placed in a determined location that improves the overall slack of the system.

BRIEF DESCRIPTION OF THE DRAWINGS

The features and advantages of the present invention are illustrated by way of example and are by no means intended to limit the scope of the present invention to the particular embodiments shown, and in which:

5 Figure 1 is a flow chart illustrating a method for designing a system according to an embodiment of the present invention;

Figure 2 illustrates a target device utilizing FPGAs according to an embodiment of the present invention;

Figure 3 illustrates a LAB according to an embodiment of the present invention;

10 Figure 4 is a flow chart illustrating a method for performing targeted logic replication for timing driven optimization according to an embodiment of the present invention;

Figure 5 illustrates an example of how criticality is determined according to an embodiment of the present invention;

15 Figure 6 is a flow chart illustrating a method for performing incremental placement according to an embodiment of the present invention;

Figure 7 illustrates fanin, fanout, and sibling relationship move proposals according to an embodiment of the present invention;

Figure 8 illustrates an exemplary critical vector move proposal according to an embodiment of the present invention;

20 Figure 9 illustrates horizontal and vertical cut-lines used for local congestion estimation according to an embodiment of the present invention;

Figure 10 is a flow chart illustrating a method for performing incremental placement utilizing directed hill-climbing according to an embodiment of the present invention;

25 Figure 11 illustrates a component trapped in a local minima according to an embodiment of the present invention;

Figure 12 illustrates basin-filling according to an embodiment of the present invention;
and

Figure 13 is a flow chart illustrating a method for managing duplicated components
according to an embodiment of the present invention.

DETAILED DESCRIPTION

Figure 1 is a flow chart that illustrates a method for designing a system according to an embodiment of the present invention. The method may be performed with the assistance of an EDA tool, for example. At 101, synthesis is performed. Synthesis includes generating a logic design of the system to be implemented by a target device. According to an embodiment of the present invention, synthesis generates an optimized logical representation of the system from a Hardware Description Language (HDL) design definition. The optimized logical representation of the system may include a representation that includes a minimized number of logic gates and logic elements required for the system. Alternatively, the optimized logical representation of the system may include a representation that has a reduced depth of logic and that generates a lower signal propagation delay.

Figure 2 illustrates an exemplary target device 200 utilizing FPGAs according to an embodiment of the present invention. The present invention may be used to design a system onto the target device 200. According to one embodiment, the target device 200 is a chip having a hierarchical structure that may take advantage of wiring locality properties of circuits formed therein. The lowest level of the hierarchy is a logic element (LE) (not shown). An LE is a small unit of logic providing efficient implementation of user logic functions. According to one embodiment of the target device 200, an LE may include a 4-input lookup table (LUT) with a configurable flip-flop.

The target device 200 includes a plurality of logic-array blocks (LABs). Each LAB is formed from 10 LEs, LE carry chains, LAB control signals, LUT chain, and register chain connection lines. LUT chain connections transfer the output of one LE's LUT to the adjacent LE for fast sequential LUT connections within the same LAB. Register chain connection lines transfer the output of one LE's register to the adjacent LE's register within a LAB. LABs are grouped into rows and columns across the target device 200. A first column of LABs is shown as 210 and a second column of LABs is shown as 211.

The target device 200 includes memory blocks (not shown). The memory blocks may be, for example, dual port random access memory (RAM) blocks that provide dedicated true dual-port, simple dual-port, or single port memory up to various bits wide at up to various frequencies. The memory blocks may be grouped into columns across the target device in between selected

5 LABs or located individually or in pairs within the target device 200.

The target device 200 includes digital signal processing (DSP) blocks (not shown). The DSP blocks may be used to implement multipliers of various configurations with add or subtract features. The DSP blocks include shift registers, multipliers, adders, and accumulators. The DSP blocks may be grouped into columns across the target device 200.

10 The target device 200 includes a plurality of input/output elements (IOEs) (not shown). Each IOE feeds an I/O pin (not shown) on the target device 200. The IOEs are located at the end of LAB rows and columns around the periphery of the target device 200. Each IOE includes a bidirectional I/O buffer and a plurality of registers for registering input, output, and output-enable signals. When used with dedicated clocks, the registers provide performance and interface

15 support with external memory devices.

The target device 200 includes LAB local interconnect lines 220-221 that transfer signals between LEs in the same LAB. The LAB local interconnect lines are driven by column and row interconnects and LE outputs within the same LAB. Neighboring LABs, memory blocks, or DSP blocks may also drive the LAB local interconnect lines 220-221 through direct link connections.

20 The target device 200 also includes a plurality of row interconnect lines ("H-type wires") 230 that span fixed distances. Dedicated row interconnect lines 230, that include H4 231, H8 232, and H24 233 interconnects, route signals to and from LABs, DSP blocks, and memory blocks within the same row. The H4 231, H8 232, and H2 233 interconnects span a distance of up to four, eight, and twenty-four LABs respectively, and are used for fast row connections in a

25 four-LAB, eight-LAB, and twenty-four-LAB region. The row interconnects 230 may drive and be driven by LABs, DSP blocks, RAM blocks, and horizontal IOEs.

The target device 200 also includes a plurality of column interconnect lines ("V-type wires") 240 that operate similarly to the row interconnect lines 230. The column interconnect lines 240 vertically routes signals to and from LABs, memory blocks, DSP blocks, and IOEs. Each column of LABs is served by a dedicated column interconnect, which vertically routes 5 signals to and from LABs, memory blocks, DSP blocks, and IOEs. These column interconnect lines 240 include V4 241, V8 242, and V16 243 interconnects that traverse a distance of four, eight, and sixteen blocks respectively, in a vertical direction.

Figure 2 illustrates an exemplary embodiment of a target device. It should be appreciated that a system may include a plurality of target devices, such as that illustrated in Figure 2, 10 cascaded together. It should also be appreciated that the target device may include programmable logic devices arranged in a manner different than that on the target device 200. A target device may also include components other than those described in reference to the target device 200. Thus, while the invention described herein may be utilized on the architecture described in Figure 2, it should be appreciated that it may also be utilized on different architectures, such as those 15 employed by Altera® Corporation in its APEX™, and Mercury™ family of chips and those employed by Xilinx®, Inc. in its Virtex™ and Virtex™ II line of chips.

Figure 3 illustrates a LAB or clustered logic block 300 according to an embodiment of the present invention. The LAB 300 may be used to implement any of the LABs shown in Figure 2. LEs 301-303 illustrate a first, second, and tenth LE in the LAB 300. The LEs 301-303 each 20 have a 4-input lookup table 311-313, respectively, and a configurable register 321-323s, respectively, connected at its output. The LAB 300 includes a set of input pins 340 and a set of output pins 350 that connect to the general-purpose routing fabric so that LAB can communicate with other LABs. The inputs to lookup tables 311-313 can connect to any one of the input pins 340 and output pins 350 using the appropriate configuration bits for each of the multiplexers 330. 25 The number of LEs, n_E , input pins, n_I , and output pins, n_O in a LAB impose important architectural constraints on a system. In addition, since a single clock line 361 and a single

asynchronous set/reset line 362 is attached to each configurable register 321-323, the configurable registers 321-323 must be clocked by the same signal and initialized by the same signal. The number of clock lines available in a LAB is represented by n_C . The number of reset lines available in a LAB is represented by n_R .

5 Referring back to Figure 1, at 102, the optimized logical design of the signal is mapped. Mapping includes determining how to implement components such as logic gates and logic elements in the optimized logic representation with specific resources on a target device. According to an embodiment of the present invention, a netlist is generated from mapping. The netlist illustrates how the resources of the target device are utilized to implement the system. The
10 netlist may, for example, include a representation of the components on the target device and how the components are connected.

At 103, the mapped logical system design is placed. Placement includes fitting the system on the target device by determining which resources on the target device are to be used for specific logic gates, logic elements, and connections between components. The placement
15 procedure may be performed by a placer in an EDA tool that utilizes placement algorithms. According to an embodiment of the present invention, a user (designer) may provide input to the placer by specifying placement constraints. The constraints may include defining logic regions that group certain components of a system together. The components may be for example, digital logic, memory devices, or other components. The size of the logic regions may be determined by
20 the user or by a sizing method. The placement of the logic regions may be determined by the user or by a placement method.

At 104, layout-driven optimizations are performed. According to an embodiment of the present invention, routing delays for the connections on the netlist are estimated by calculating a fastest possible route. Timing-driven netlist optimization techniques may be applied to perturb
25 the netlist to reduce the critical path(s). The netlist may be perturbed by the EDA tool performing the synthesis, mapping and placement. Alternatively, the netlist may be perturbed by a user of

the EDA tool, or by a third party. Perturbing the netlist may include adding, deleting, or moving components. According to an embodiment of the present invention, preferred locations are identified for the components that have been added or moved from the layout-driven optimization. The locations assigned to components of the existing system from the placement 5 procedure are identified as preferred locations for the components.

According to an embodiment of the present invention, layout-driven optimizations includes duplicating a component having a specific LUT and/or register that results in reducing a length of a critical path without significantly affecting the lengths of other paths in the circuit.

The system design is altered such that a component on the critical path is copied and relocated to 10 reduce the delay along the critical path. The delay reduction on the critical path results in an increase of the maximum frequency at which the system can operate.

Figure 4 is a flow chart illustrating a method for performing targeted logic replication for timing driven optimization according to an embodiment of the present invention. According to one embodiment, each atom on a netlist corresponding to a resource on a PLD is examined to 15 determine its criticality using the procedure described.

At 401, the criticality of a component in a system is determined. The criticality of a component may be determined by evaluating the slack on a connection associated with the component. The criticality may be quantified, for example, by the ratio of the slack of the connection to the maximum permissible delay of the connection. Alternatively, criticality may be 20 quantified by the slack value associated with the connection alone, a ratio of the connection's delay and the maximum permissible delay of the connection, or other metrics. A positive slack value for a connection indicates how much delay on the connection in isolation can change before a worst-case path through that connection just meets its timing constraints. A negative slack value for a connection indicates how much the delay on the connection needs to change in order 25 for that connection to meet its timing constraint. According to an embodiment of the present invention, a slack value for a connection may be determined by computing a cumulative delay for

the path. The cumulative delay for the path may be computed by summing the estimated delays (achievable delays) for transmitting a signal between the functional blocks on the path. The cumulative delay for the path is subtracted from the long-path timing constraint of the path to determine a path slack value. The smallest path slack value is the slack value for the connection.

5 Figure 5 illustrates an example of how slack values are generated. In this example, a first path passes from register R1 510 to functional block FB1 511 to register R2 512. A second path passes from register R3 513 to functional block FB1 511 to register R2 512. A long-path timing constraint of 10 ns is specified for the first path from R1 510 to R2 512. A long-path timing constraint of 20 ns is specified for the second path from R3 513 to R2 512. The long-path timing
10 constraint may be associated with a FMAX (maximum clocking frequency), TSU (setup time), TCO (clock-to-output time), or other specification.

To determine a slack value for a connection, for each path through the connection, a cumulative delay for the path is computed. This may be computed by summing the achievable delays associated with the connections on the path. The achievable delay for each connection is
15 shown next to the connection. For the first path through R1 510, FB1 511, and R2 512, the cumulative delay is 5 ns. The cumulative delay, 5 ns, is subtracted from the long-path timing constraint for the first path, 10 ns, to determine a path slack of 5ns. For the second path through R3 513, FB1 511, and R2 512, the cumulative delay is 3.5 ns. The cumulative delay, 3.5 ns, is subtracted from the long-path timing constraint for the path, 20 ns, to determine a path slack
20 value 16.5 ns. The path slack value having the lowest value is the slack value for the connection. For connection 521 between R1 510 and FB1 511, the path slack value is 5 ns. For connection 522 between FB1 511 and R2 512, the path slack value is 5 ns. For connection 523 between R3 513 and FB1 511, the path slack value is 16.5 ns. It should be appreciated that other types of procedures and techniques may also be used to access the criticality of a component.

25 Referring back to Figure 4, at 402, it is determined whether the criticality exceeds a first predetermined threshold value. If the criticality does not exceed the first predetermined threshold

value, control proceeds to 403. If the criticality exceeds the first predetermined threshold value, control proceeds to 404.

At 403, the component is not duplicated.

At 404, it is determined whether the path delay of the connection associated with the

5 component exceeds a second predetermined threshold value. Path delay may be described as the time required for transmitting a signal on the connection given the current configuration of the system. If the path delay of the connection exceeds the second predetermined threshold value, control returns to 403. If the path delay of the connection does not exceed the second predetermined threshold value, control proceeds to 405.

10 At 405, the component is designated as a replication candidate.

At 406, a location for placing the duplicate of the replication candidate is determined.

According to an embodiment of the present invention, all locations on the target device are considered for a location to place the duplicate. According to an alternate embodiment of the present invention, a search space on the target device is defined to narrow the considered

15 locations. In this embodiment, the search space may be an area proximate to the destination of the critical signal associated with the replication candidate such as a location within the same LAB as the destination of the critical signal. According to an embodiment of the present invention, the location for placing the duplicate is selected without considering whether or not placement at a particular location would result in architectural violations or illegalities.

20 At 407, the duplicate of the replication candidate is assigned the connection from the replication candidate.

At 408, the criticality of the duplicate of the replication candidate is determined. The criticality of a component may be determined by evaluating the slack on the connections associated with the component. As described in 401, the criticality may be quantified, for

25 example, by the ratio of the slack of the connection to the maximum permissible delay of the connection. Alternatively, criticality may be quantified by the slack value associated with the

connection alone, a ratio of the connection's delay and the maximum permissible delay of the connection, or other metrics. In a preferred embodiment of the present invention, the slack values associated with the connections surrounding the duplicate of the replication candidate are compared with the slack values associated with the connections surrounding the component.

5 At 409, the slack changes of connections surrounding the duplicate replication candidate are compared with the original slack of the connections surrounding the component determined at 401. The comparison determines whether the replication yields an improvement or a positive "slack gain". If the slack gain is negative, control proceeds to 403. If the slack gain is positive, control proceeds to 410.

10 At 410, gain values are generated for the replication candidate. According to an embodiment of the present invention, gain values are generated in response to the slack gain determined at 409, the criticality of the move from the slack associated with the replication candidate determined at 401, and illegalities associated with placement of the duplicate at the location determined at 406. According to an embodiment of the present invention, illegalities 15 associated with placement of the duplicate may include a high level of congestion at the placement location. According to an embodiment of the present invention, after each atom on the netlist is evaluated using the procedure described with reference to Figure 4, the gain values generated for the replication candidates are utilized to determine which replication candidates are to be selected for replication. In one embodiment, the replication candidates with the n highest 20 gain values are selected to be replicated. It should be appreciated that n may be any value.

According to an embodiment of the present invention, control proceeds to determine whether an additional connection (fanout) is produced by the component at 401. If no additional fanouts exist, the component is deleted. If an additional connection is driven by the component, the connection is moved to the replicated candidate and the slack on the connection at 25 its original location is compared with the slack on the connection at its new location. If there is a slack gain, the connection is assigned to the replicated candidate.

Referring back to Figure 1, at 105, incremental placement is performed on duplicates of replicated candidates placed at locations that result in architectural violations or illegalities. The changes to the netlist generated from layout-driven optimization are placed on the layout of the existing system placed at 103. Incremental placement involves evaluating resources on a target device such as LABs that have architectural violations or illegalities from layout-driven optimizations. Incremental placement attempts to perturb the preferred locations as little as possible to ensure that the final placement respects all architectural constraints. Incremental placement attempts to identify non-critical LEs that may be moved from their preferred locations to resolve architectural violations in order that truly critical elements may stay at their preferred locations. Incremental placement may be performed by an incremental placement engine (not shown) in the EDA tool that utilizes incremental placement algorithms.

In performing incremental placement, an architectural description of the target device, A, and a netlist, N(E,C), that includes a set of logic elements, E, and a set of connections, C, is processed. Each element, e, is associated with a preferred physical location, ($p_x(e)$, $p_y(e)$).
According to an embodiment of the present invention, all atoms of the netlist have a preferred location. Incremental placement generates a set of mapped locations, M, for each logic elements in N. Incremental placement tries to find a mapping from preferred locations to mapped locations, $P \rightarrow M$, such that the mapped locations are architecturally feasible as well as being minimally disruptive. The definition of minimal disruption depends on the goal of netlist optimization.

According to an embodiment of the present invention, the goal of netlist optimization is to optimize timing of the system. In this embodiment, T(S) represent an estimate of the critical path delay if all logic elements in E are mapped to ($s_x(e)$, $s_y(e)$). The estimate may ignore the legality of locations and may be computed assuming a best case route is possible for each connection. In this example, $P \rightarrow M$ is minimally disruptive if incremental placement minimizes { $T(M)-T(P)$ }. Any logic element can be moved from its preferred location as long as it does not

degrade the critical path. According to one embodiment, routing area is also tracked to control excessive routing congestion. In this embodiment, $A(S)$ represents the routing area consumed if the logic elements are mapped to $(s_x(e), s_y(e))$. Minimal disruptiveness is satisfied by minimizing the relationships shown below.

$$5 \quad \{T(M)-T(P)\} + \{A(M)-A(P)\} \quad (1)$$

Figure 6 is a flow chart illustrating a method for performing incremental placement according to an embodiment of the present invention. The method described in Figure 6 may be used to perform incremental placement as shown as 105 in Figure 1. At 601 proposed moves for all LEs in a LAB having architectural violations are generated. According to an embodiment of the present invention, proposed moves may include a move-to-fanin, move-to-fanout, move-to-sibling, move-to-neighbor, move-to-space, a move towards a critical vector, and other moves. A move-to-fanin involves moving an LE to a LAB that is a fanin of the LE. A move-to-fanout involves moving an LE to a LAB that is a fanout of the LE. A move-to-sibling involves moving an LE to a LAB that is fanout of a LAB that fans in to the LAB of the LE.

15 Figure 7 illustrates examples of a move-to-fanin, move-to-fanout, and move-to-sibling.
When a first LE in a first LAB transmits a signal to a second LE in a second LAB, the first LAB
is said to be a fanin of the second LE. When a first LE in a first LAB receives a signal from a
second LE in a second LAB, the first LAB is said to be a fanout of the second LE. When a first
LE from a first LAB receives a signal from a second LE from a second LAB that also transmits to
20 a third LE in a third LAB, the first LAB and the third LABs are said to be siblings. Blocks 701-
709 illustrate a plurality of LABs. Each of the LABs 701-709 has a number of shown LEs. A
plurality of arrows 711-718 are shown to illustrate the direction of a signal transmitted between
LEs. Relative to LAB 706, LABs 701-704 are considered fanins, LABs 705 and 707 are
considered siblings, and LABs 708 and 709 are considered fanouts.

25 Proposed moves may also include move-to-neighbor, move-to-space, and move towards
critical vector. A move-to-neighbor involves moving an LE to an adjacent LAB. A move-to-

space involves a move to any random free LE location in a target device. A move towards critical vector involves moving an LE towards a vector that is computed by summing the directions of all critical connections associated with the moving LE. Figure 8 illustrates an exemplary critical vector 801. Vector 801 is the critical vector of LE 811 which has critical 5 connections to LEs 812 and 813, and a non-critical connection with LE 814.

Referring back to Figure 6, at 602, a current placement of LEs in a LAB with architectural violations and proposed moves of the LEs in the LAB are evaluated by a cost function. The cost function may include parameters which measure the legality of a LAB (cluster 10 legality cost), timing (timing cost), and an amount of routing resources that is required for a placement (wirelength cost). According to an embodiment of the present invention, the cost function guides the reduction of architectural violations while ensuring minimal disruption. This cost function, C, is illustrated with the relationship shown below.

$$C = K_L * \text{ClusterCost} + K_T * \text{TimingCost} * K_W * \text{WirelengthCost} \quad (2)$$

K_L , K_T , and K_W represent weighting coefficients that normalize the contributions of each 15 parameter. It should be appreciated that other parameters may be used in addition to or in place of the parameters described.

The cluster legality cost is a cost associated with each LAB CL_i . This cost may be represented as shown below.

$$\text{ClusterCost}(CL_i) = kE_i * \text{legality}(CL_i, n_E) + \\ KI_i * \text{legality}(CL_i, n_I) + \\ kR_i * \text{legality}(CL_i, n_R) + \\ kO_i * \text{legality}(CL_i, n_O) + \\ kC_i * \text{legality}(CL_i, n_C) \quad (3)$$

The legality (CL_i, \dots) function returns a measure of legality for a particular constraint. A 25 value of 0 indicates legality, while any positive value is proportional to the amount to which the constraint has been violated. Functions $\text{legality}(CL_i, n_E)$, $\text{legality}(CL_i, n_I)$, $\text{legality}(CL_i, n_O)$, $\text{legality}(CL_i, n_R)$, and $\text{legality}(CL_i, n_C)$ evaluate if LAB CL_i has a feasible number of logic elements, inputs, outputs, reset lines and clock lines, respectively. According to an embodiment

of the present invention, the weighting coefficients kE_i , KI_i , kO_i , kR_i , and kC_i are all initially set to 1 for every LAB CL_i in the target device.

The timing cost associated with a placement may be represented as shown below.

$$\text{TimingCost} = \text{TC}_{\text{VPR}} + k_{\text{DAMP}} * \text{TC}_{\text{DAMP}} \quad (4)$$

5 The first parameter, TC_{VPR} , is based upon the cost used by a versatile placement and routing (VPR) placer. This cost may be represented with the following relationship.

$$\text{TC}_{\text{VPR}} = \sum_c \text{crit}(c) * \text{delay}(c) \quad (5)$$

This function encourages critical connections to reduce delay while allowing non-critical connections to optimize wirelength and other optimization criteria.

10 The second parameter, TC_{DAMP} , operates as a damping component of the timing cost function and can be represented with the following relationships.

$$\text{TC}_{\text{DAMP}} = \sum_c \max(\text{delay}(c) - \text{maxdelay}(c), 0.0) \quad (6)$$

$$\text{maxdelay}(c) = \text{delay}(c) + \alpha * \text{slack}(c) \quad (7)$$

15 The damping component penalizes any connection c whose $\text{delay}(c)$ exceeds a maximum value $\text{maxdelay}(c)$. This allows arbitrary moves to be made along a plateau defined by the maximum delays. The maxdelay values may be updated every time a timing analysis of the system is executed. The maxdelay values are controlled by the slack on the connection considered. The parameter α determines how much of a connection's slack will be allocated to the delay growth of the connection. Thus, the plateau is defined by the connection slack so that
20 connection with large amounts of slack are free to move large distances in order to resolve architectural violations, while small slack values are relatively confined.

Wirelength cost of a placement may be measured by determining a number of routing wires that cross cut-lines that outline a LAB. Figure 9 illustrates the utilization of cut-lines according to an embodiment of the present invention. Blocks 901-909 represent LABs having a
25 plurality of shown LEs. Horizontal cut-lines 911 and 912 and vertical cut-lines 913 and 914 are placed in each horizontal channel of a target device. Cut-lines provide a method to measure

congestion by finding the regions that have the largest number of routing wires 921-924. This measurement may be used to prevent the formation of localized congested areas that can cause circuitous routes. The total number of routing wires that intersect a particular cut may be calculated by finding all the signals that intersect a particular cut-line and summing the average 5 crossing-count for each of these signal wires. The average crossing count for a signal may be computed using the following relationship.

$$\text{CrossingCount(net)} = q(\text{NumCLBlockPins(net)}) \quad (8)$$

The function q is given as a number of discrete crossing counts as a function of signal pin count. The argument to the function q is the number of clustered logic block pins used to wire the 10 signal. With respect to the functions shown in (3)-(8), it should be appreciated that other types of functions may be used in addition or in place of the functions represented.

Referring back to Figure 6, at 603, it is determined whether the cost associated with any of the proposed moves is better than the cost associated with the current placement. The costs associated with the proposed moves and current placement may be obtained by using cost 15 function values generated from using the cost function described with respect to 602. If it is determined that the cost associated with any of the proposed moves is better than the cost associated with the current placement, control proceeds to 604. If it is determined that the cost associated with any of the proposed moves is not better than the cost associated with the current placement, control proceeds to 605.

At 604, the proposed move associated with the best cost is selected as the current 20 placement.

At 605, it is determined whether any additional LABs in the system have architectural violations. If additional LABs in the system have architectural violations, control will move to one of these LABs and proceeds to 601. If no additional LABs in the system have architectural 25 violations, control proceeds to 606 and terminates the procedure. According to an embodiment of the present invention, a counter may be used to track the number of proposed moves that have

been generated, or the number of LEs or LABs that have had proposed moves generated. In this embodiment, when this number exceeds a threshold value, instead of proceeding to 601, control terminates the procedure and returns an indication that a fit was not found.

Figure 10 is a flow chart illustrating a method for performing incremental placement

5 utilizing directed hill-climbing according to an embodiment of the present invention. The method described in Figure 10 may be used to perform incremental placement as shown as 105 in Figure

1. At 1000, a loop iteration index, L, is set to 1.

At 1001 proposed moves for all LEs in a LAB having architectural violations are generated. According to an embodiment of the present invention, the proposed moves may be

10 generated similarly as described in 601 shown in Figure 6. The number of LEs having proposed moves generated is recorded.

At 1002, a current placement of LEs in a LAB with architectural violations and proposed moves of the LEs in the LAB are evaluated by a cost function. According to an embodiment of the present invention, the evaluation performed may be similarly conducted as described in 602

15 of Figure 6.

At 1003, it is determined whether the cost associated with any of the proposed moves is better than the cost associated with the current placement. The costs associated with the proposed moves and current placement may be obtained by using values generated from using the cost function described with respect to 602. If the cost associated with any of the proposed moves is

20 better than the cost associated with the current placement, control proceeds to 1004. If the cost associated with any of the proposed moves is not better than the cost associated with the current placement, control proceeds to 1005.

At 1004, the proposed move associated with the best cost is selected as the current placement.

25 At 1005, it is determined whether any additional LABs in the system have architectural violations. If additional LABs in the system have architectural violations, control will move to

one of these LABs and proceeds to 1007. If no additional LABs in the system have architectural violations, control proceeds to 1006 and terminates the procedure.

At 1007, it is determined whether the number of LEs that have proposed moves generated exceeds the value K where K is a predefined value. If the number of LEs that have proposed moves generated exceeds the value K, control proceeds to 1009. If the number of LEs that have proposed moves generated does not exceed the value K, control proceeds to 1008.

At 1008, the loop iteration index, L, is incremented. Control returns to 1001.

At 1009, timing analysis is performed. According to an embodiment of the present invention, the values for maxdelay and crit(c), used for evaluating timing cost, are updated to reflect the current configuration of the system.

At 1010, the cost function is updated. According to an embodiment of the present invention, weighting coefficients in the ClusterCost parameter are incremented in proportion to an amount of violation. Updating the cost function allows directed hill-climbing to be performed. Directed hill-climbing is a technique that is used for generating proposed moves when moves cannot be found to decrease the current cost of a placement.

Figure 11 illustrates an example where directed hill-climbing may be applied. The target device 1100 includes a plurality of LABs 1101-1105 each having a plurality of shown LEs. In this example, LAB 1103 has one LE more than is allowed by its architectural specification.

Every possible move attempt to resolve the architectural constraints of the center LAB 1103 results in another architectural violation. If all architectural violations are costed in the same manner, then the method described in Figure 6 may have difficulties resolving the constraint violation.

Figure 12 illustrates a two dimensional slice of the multi-dimensional cost function described. The current state 1201 represents the situation shown in Figure 11. No single move in the neighborhood of the current state finds a solution with a lower cost. However, the cost function itself could be modified to allow for the current state 1201 to climb the hill. The

weighting coefficients of the cost function may be gradually increased for LABs that have unsatisfied constraints. A higher weight may be assigned to unsatisfied constraints that have been violated over a long period of time or over many iterations. This results in the cost function being reshaped to allow for hill climbing. The reshaping of the cost function has the effect of filling a 5 basin where the local minima is trapped. Referring back to Figure 11, once the weighting coefficients have been increased for LAB 1103, a proposed move to one of the adjacent cluster may be made to allow for shifting the violation "outwards" to a free space.

Updating a cost function also allows for a quick convergence by preventing a phenomenon known as thrashing. Thrashing occurs when incremental placement is trapped in an 10 endless cycle where an LE is moved between two points in the configuration space which both result in architectural violations. By increasing the cost or penalty for moving to the two points, a move to a third point would eventually be more desirable and accepted.

Referring back to Figure 10, at 1011, it is determined whether the loop index, L, is greater than a threshold value. If the loop index, L, is not greater than the threshold value, control 15 proceeds to 1008. If the loop index, L, is greater than the threshold value, control proceeds to 1012.

At 1012, control terminates the procedure and returns an indication that a fit was not found.

Referring back to Figure 1, at 106 it is determined whether additional optimization is to 20 be performed. According to an embodiment of the present invention, a count is maintained as to how many times the procedure at 104 is performed. If the count does not exceed a predefined threshold number, control returns to 104. If the count exceeds the predefined threshold number, control proceeds to 107.

According to an embodiment of the present invention, additional procedures may be 25 performed on system. Figure 13 is a flow chart illustrating a method for managing duplicated components according to an embodiment of the present invention. At 1301, it is determined

whether a duplicate has been placed outside a predefined range from its corresponding replication candidate. According to one embodiment, the predefined range may be, for example, a distance of a logic-array block. If the duplicate is within the range, control proceeds to 1302. If the duplicate is place outside the range, control proceeds to 1303.

5 At 1302, the duplicate is removed and the connection associated with the duplicate is restored to the replication candidate.

At 1303, the system slack for the current design is determined. According to an embodiment of the present invention, the system slack for the current design reflects the overall slack of the current design.

10 At 1304, it is determined whether the system slack for the current design is an improvement over the previous design. If the system slack for the current design is not an improvement over the previous design, control proceeds to 1305. If the system slack for the current design is an improvement over the previous design, control terminates the procedure as shown at 1306.

15 At 1305, the previous design for the system is restored. According to an embodiment of the present invention, the netlist and placement for each design is saved before performing the procedure at 104 to allow the previous design to be restored.

Referring back to Figure 1, at 107, routing of the system is performed. During routing, routing resources on the target device are allocated to provide interconnections between logic gates, logic elements, and other components on the target device. The routing procedure may be performed by a router in an EDA tool that utilizes routing algorithms. According to an alternate embodiment of the present invention, it should be appreciated that the procedures with reference to 104-106 may also be performed after the routing described with reference to 107.

The incremental placement techniques disclosed allow logic changes to be incorporated 25 into an existing system design without reworking placement of the entire system. The incremental placement techniques attempt to minimize disruption to the original placement and

maintain the original timing characteristics. According to an embodiment of the present invention, a method for designing a system on a target device utilizing FPGAs is disclosed. The method includes placing new LEs at preferred locations on a layout of an existing system. Illegibilities in placement of the components are resolved. According to one embodiment, 5 resolving the illegalities in placement may be achieved by generating proposed moves for an LE, generating cost function values for a current placement of the LE and for placements associated with the proposed moves, and accepting a proposed move if its associated cost function value is better than the cost function value for the current placement.

Figures 1, 6, 10 and 13 are flow charts illustrating methods according to embodiments of 10 the present invention. Some of the techniques illustrated in these figures may be performed sequentially, in parallel or in an order other than that which is described. It should be appreciated that not all of the techniques described are required to be performed, that additional techniques may be added, and that some of the illustrated techniques may be substituted with other techniques.

15 Embodiments of the present invention (e.g. exemplary process described with respect to Figures 1, 4, and 5) may be provided as a computer program product, or software, that may include a machine-readable medium having stored thereon instructions. The machine-readable medium may be used to program a computer system or other electronic device. The machine-readable medium may include, but is not limited to, floppy diskettes, optical disks, CD-ROMs, 20 and magneto-optical disks, ROMs, RAMs, EPROMs, EEPROMs, magnetic or optical cards, flash memory, or other type of media/machine-readable medium suitable for storing electronic instructions.

In the foregoing specification the invention has been described with reference to specific 25 exemplary embodiments thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the

invention. The specification and drawings are, accordingly, to be regarded in an illustrative rather than restrictive sense.